



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/001,317	10/23/2001	Michael Kowalchik	EMR-00401	2338
25181	7590	06/23/2004	EXAMINER	
FOLEY HOAG, LLP PATENT GROUP, WORLD TRADE CENTER WEST 155 SEAPORT BLVD BOSTON, MA 02110			CHACE, CHRISTIAN	
			ART UNIT	PAPER NUMBER
			2187	
DATE MAILED: 06/23/2004				

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/001,317	KOWALCHIK ET AL.	
	Examiner Christian P. Chace	Art Unit 2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 08 June 2004.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-23,25,30,32 and 34-42 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-23,25,30,32 and 34-42 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 23 October 2001 is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892)	4) <input type="checkbox"/> Interview Summary (PTO-413)
2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)	Paper No(s)/Mail Date. _____
3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.	5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152)
	6) <input type="checkbox"/> Other: _____.

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 8 June 2004 has been entered.

Response to Amendment

Previously un-entered amendment filed 12 May 2004 has now been entered. Claims 1-23, 25, 30, 32, and 34-42 are pending. This action is NOT final.

Drawings

The drawings are objected to under 37 CFR 1.83(a). The drawings must show every feature of the invention specified in the claims. Therefore, the storage in response to a power failure must be shown or the feature(s) canceled from the claim(s). No new matter should be entered.

Corrected drawing sheets are required in reply to the Office action to avoid abandonment of the application. Any amended replacement drawing sheet should include all of the figures appearing on the immediate prior version of the sheet, even if only one figure is being amended. The figure or figure number of an amended drawing should not be labeled as "amended." If a drawing figure is to be canceled, the appropriate figure must be removed from the replacement sheet, and where necessary,

the remaining figures must be renumbered and appropriate changes made to the brief description of the several views of the drawings for consistency. Additional replacement sheets may be necessary to show the renumbering of the remaining figures. The replacement sheet(s) should be labeled "Replacement Sheet" in the page header (as per 37 CFR 1.84(c)) so as not to obstruct any portion of the drawing figures. If the changes are not accepted by the examiner, the applicant will be notified and informed of any required corrective action in the next Office action. The objection to the drawings will not be held in abeyance.

Double Patenting

Copending applications 10/001,317 (instant) and 10/004,090 were found to have obviousness-type provisional double-patenting issues. 10/004,090 was found to have broader claims than the instant application.

The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. See *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and, *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the conflicting application or patent is shown to be commonly owned with this application. See 37 CFR 1.130(b).

Effective January 1, 1994, a registered attorney or agent of record may sign a terminal disclaimer. A terminal disclaimer signed by the assignee must fully comply with 37 CFR 3.73(b).

Claims 4, 5, and 23 are provisionally rejected under the judicially created doctrine of obviousness-type double patenting as being unpatentable over claims 1, 4, and 20,

Art Unit: 2187

respectively, of copending Application No. 10/004,090. Although the conflicting claims are not identical, they are not patentably distinct from each other because claims 4 and 5 of the instant application are obvious over respective claims 1 and 4 of the copending application; and instant claim 23 is missing one limitation of copending application's claim 20, which is an obvious modification. All of this is explained below in detail.

This is a provisional obviousness-type double patenting rejection because the conflicting claims have not in fact been patented.

With respect to claims 4 and 5 of the instant application, please refer to the table below, which illustrates the obvious and anticipatory relationship of the claim limitations at issue:

<u>Instant Application 10/001,317</u>	<u>Application 10/004,090</u>
4. A cache comprising:	1. A data storage device, the device comprising:
A front-end interface	A device interface
That receives data access requests	For receiving data access requests
That specify respective data storage addresses	[inherent: any storage access request must specify an address, or location, of the data or area it wishes to access]
A back-end interface that can retrieve data identified by the data storage address	
Cache storage formed by at least two disks	More than two disk drives
Cache manager	A controller
That services at least some of the requests	That accesses the disk drives
Received at the front-end interface	In response to the received data access requests
Using data stored in the cache storage	[inherent: by accessing the disk drives, data stored in the disk drives is "used."]
Where said cache manager stores, in response to detecting a power failure, identification data identifying addresses within the cache storage where data is	Obvious – see <i>Graham v. Deere</i> analysis below, in view of Katz et al (US Patent

stored and the corresponding addresses at a back-end storage area where the data is stored	#5,195,100).
Wherein the disks comprise disks having platters less than 3.5 inches in diameter	[disk drives] having platter sizes less than 3.5 inches in diameter
5. (depends upon claim 4, so contains all of the limitations supra)	4. (depends on claim 1, so contains all of the limitations supra)
Wherein the disks comprise disks having at least one of the following platter sizes: 2.5 inches, 1.8 inches, and 1 inch in diameter	Wherein the platter sizes comprise platters of at least one of the following sizes: 2.5 inches, 1.8 inches, and 1 inch

The cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy.

The difference between instant claims 4 and 5 and copending claims 1 and 4 is performing such storage in response to a power failure.

Katz et al discloses storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information as a result of the power failure, as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

Claim 23 is provisionally rejected under the judicially created doctrine of obviousness-type double patenting as being unpatentable over claim 20 of copending Application No. 10/004090 in view of *Modern Operating Systems*, Tanenbaum, page 306 (cited in IDS filed 10/23/01 in both applications). As will be illustrated below in the table, the instant claim 23 would anticipate copending claim 20, with the exception of the at least one first data storage device having a platter size of at least 3.5 inches in diameter.

<u>Instant Application 10/001,317</u>	<u>Application 10/004,090</u>
23. A data storage system, comprising:	20. A data storage system, the system comprising:
A back-end storage system having an address space, addresses in the address space identifying blocks of storage	At least one first data storage device [inherent: storage systems must have addresses and addresses inherently identify blocks of storage, or storage locations]
[Obvious: see <i>Graham v. Deer</i> factors below]	Having a platter size of at least 3.5 inches in diameter
A cache for the back-end storage system Having a lesser storage capacity than the back-end system	At least one second data storage device [inherent: definition of a cache, i.e., smaller and faster memory for faster access closer to the processor]
The cache including:	Comprising:
A front-end interface	A device interface
That receives I/O requests	For receiving data access requests
That specify respective addresses of back-end storage blocks	[inherent as discussed <i>supra</i>]
A back-end interface	A second controller
That communicates with the back-end storage system	That coordinates data access to the at least one first data storage device and the at least one second storage device
Cache storage formed by at least two disks having platter diameters less than 3.5 inches	More than two disk drives coupled to the controller, the drivers having platter sizes less than 3.5 inches in diameter
A cache manager that services at least some of the I/O requests received via the front-end interface using blocks temporarily stored in the cache storage	A first controller configured to receive data access requests from the interface

Where said cache manager further stores, in response to detecting a power failure, identification data identifying addresses of data stored within said cache storage and the corresponding addresses at said back-end storage area where the data is stored	Obvious, see <i>Graham v. Deere</i> analysis below, in view of Katz et al (US Patent #5,195,100).
--	---

However, Tanenbaum discloses that CD ROM's are 120 millimeters across, which is more than 3.5 inches, in the second paragraph of page 306.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the instant application and Tanenbaum before him/her, to have a back-end storage device (first data storage device) with platter sizes of at least 3.5 inches, as CD ROMS are at least 3.5 inches, and they have much higher recording densities than conventional magnetic disks, as disclosed by Tanenbaum in line 2 of the second paragraph on page 306.

The cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy.

The difference between instant claims 4 and 5 and copending claims 1 and 4 is performing such storage in response to a power failure.

Katz et al discloses storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information as a result of the power failure, as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

This is a provisional obviousness-type double patenting rejection.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claims 32-37 are rejected under 35 U.S.C. 102(b) as being anticipated by Brant et al (US Patent #5,805,787).

With respect to independent claim 32, a data storage system is disclosed in figure 1.

A 'back-end' storage system is disclosed in figure 1, #25. Data storage spaces all have addresses, or locations for storing data – this is inherent, by definition of an address and by definition of data storage. The addresses identifying "blocks" of storage is also inherent – a block of storage may be any size, as applicants have not limited such size in the instant claim. Therefore, examiner interprets a "block" to be one memory location.

A plurality of caches is disclosed in column 5, lines 29-31, which discusses that the invention of Brant et al may be put in any multiple locations in the hierarchy of the system. Inherently, if the Brant et al invention is placed in “back-to-back” locations, if you will, then the front end interface of one of the locations will connect to the back-end interface of the next. Indeed, this is the definition of a hierarchical memory scheme. Inherently, each level stores less than the main memory, or back-end storage system, by definition. The back-end storage system is disclosed in figure 1 as #16 as well as in the title and abstract. A cache is, by definition, a memory subsystem in which frequently used values are duplicated for quick access. A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory. A cache is useful between a processor and a main memory when main memory accesses are slow compared to the processor speed, because the cache memory is always faster than the main memory.

A front-end interface is disclosed in figure 1, #11, as the “host interface” in column 6, line 17. Receiving I/O requests that specify respective addresses of back-end storage blocks is disclosed in column 5, line 35, which refers to I/O rates. An I/O rate must be calculated from I/O requests. Therefore, I/O requests are inherent in a system with an I/O rate being measured. Also, it is important to note that I/O is “Input/Output,” which is defined as the complementary tasks of gathering data for a computer or program to work with, and of making the results of the computer’s activities available to the user or to other computer processes. This is exactly what figure 1 shows – data is gathered from the disk-based disk cache and supplied to the host.

A back-end interface is disclosed in figure 1 as #24, as the "interface to the mass data storage subsystem" in column 6, line 18.

Cache storage being formed by at least two disks is disclosed in figure 1, #16, and discussed in the abstract, for example. The disks having platter diameters [of] less than 3.5 inches is disclosed in column 3, lines 45-46.

A cache manager is disclosed as the storage controller in figure 1, #20. As discussed in column 4, lines 9-19, the storage controller receives data access requests (via front-end interface #11) that specify respective data storage addresses, retrieves data identified by the data storage addresses, and services at least one of the I/O requests received at the front-end interface using data stored in the cache storage. The at least one of the requests serviced are, of course, the ones in which the cache holds the information for. As the definition of cache suggests, as discussed *supra*, if the cache does not have the information, the cache does not service the request *per se*, but the main memory, or larger storage does (via back-end interface #24). This is also discussed in column 4, lines 9-19. As the data stored in the at least two disks that make up the cache is stored striped in a RAID configuration, as discussed in column 5, lines 34, 36, and 44, for example. This is significant as RAID 3, for example, as mentioned in column 5, line 44, adds redundant information in the form of parity to a parallel access striped array, by definition. Striping is assigning blocks of storage in regular sequence to all of an array's disks. Therefore, Brant et al uses stores the data in the at least two disks that make up the cache by striping it in a RAID 3 configuration, for example.

The caches having respective cache storage address space is inherent – without addresses for the cache space, the system will not know here to go for requested information, or know where to put information.

With respect to claim 34, the plurality of caches being connected in series such that the front-end interface of one of said plurality of caches is coupled to the back-end interface of another of said plurality of caches is discussed supra with respect to claim 32. A plurality of caches is disclosed in column 5, lines 29-31, which discusses that the invention of Brant et al may be put in any multiple locations in the hierarchy of the system. Inherently, if the Brant et al invention is placed in “back-to-back” locations, if you will, then the front end interface of one of the locations will connect to the back-end interface of the next. Indeed, this is the definition of a hierarchical memory scheme.

Also inherent in the hierarchical scheme discussed supra, the front-end interface of the highest level cache would inherently connect to the host, while the back-end interface of the lowest-level cache would have to connect to the back-end storage system. This again, is by definition of hierarchical memory system.

With respect to claim 35, whereupon receiving one of [the] I/O requests at said front-end interface of one of the plurality of caches, the cache manager of said one of said plurality of caches sends data corresponding to said one of [the] I/O requests to said device making said one of [the] I/O requests if data is stored on the cache storage of said one of [said] plurality of caches is inherent, as it is the definition of a cache. All the instant claim recites is that if the requested data is in the cache, the cache sends the data to the host.

With respect to claim 36, whereupon receiving one of [the] I/O requests at said front-end interface of said one of said plurality of caches, the cache manager of said one of [said] plurality of caches sends said one of [said] I/O requests to one of said back-end storage [device] and another of said plurality of caches coupled to the back-end interface of said one of [said] plurality of caches if data is not stored on [in] the cache storage of said one of said plurality of caches is inherent in a hierarchical memory system, as discussed supra. The instant claim merely recites going to the next level cache if the data is not found in the first-level cache. This is how a hierarchical system such as the one discussed supra must operate.

With respect to claim 37, the cache further comprising at least one interface conforming to a protocol to allow at least one additional disk to be connected to said storage is disclosed in column 6, lines 39-41, for example, where the protocol is a SCSI-type connection(s). A SCSI interface is inherently used to connect CPU's or hosts to SCSI peripheral devices (note the plural), such as many hard disks and printers, using SCSI ports, which inherently provide logical connections between the computer and all of the devices on the SCSI bus. Therefore, inherently, a SCSI bus connection through SCSI ports allows at least one additional disk to be connected through the interface.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-23, 25, 30, and 38-42 are rejected under 35 U.S.C. 103(a) as being unpatentable over Brant et al (cited *supra*) in view of Katz et al (US Patent #5,195,100).

With respect to independent claim 1, a cache is disclosed in the title and abstract (#16 in figure 1). A cache is, by definition, a memory subsystem in which frequently used values are duplicated for quick access. A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory. A cache is useful between a processor and a main memory when main memory accesses are slow compared to the processor speed, because the cache memory is always faster than the main memory.

A front-end interface is disclosed in figure 1, #11, as the "host interface" in column 6, line 17.

A back-end interface is disclosed in figure 1 as #24, as the "interface to the mass data storage subsystem" in column 6, line 18.

Cache storage being formed by at least two disks is disclosed in figure 1, #16, and discussed in the abstract, for example.

A cache manager is disclosed as the storage controller in figure 1, #20. As discussed in column 4, lines 9-19, the storage controller receives data access requests (via front-end interface #11) that specify respective data storage addresses, retrieves data identified by the data storage addresses, and services at least one of the requests received at the front-end interface using data stored said at least two disks. The at least one of the requests serviced are, of course, the ones in which the cache holds the information for. As the definition of cache suggests, as discussed *supra*, if the cache

does not have the information, the cache does not service the request *per se*, but the main memory, or larger storage does (via back-end interface #24). This is also discussed in column 4, lines 9-19. As the data stored in the at least two disks that make up the cache is stored striped in a RAID configuration, as discussed in column 5, lines 34, 36, and 44, for example. This is significant as RAID 3, for example, as mentioned in column 5, line 44, adds redundant information in the form of parity to a parallel access striped array, by definition. Striping is assigning blocks of storage in regular sequence to all of an array's disks. Therefore, Brant et al uses stores the data in the at least two disks that make up the cache by striping it in a RAID 3 configuration, for example.

The cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy, as discussed in column 7, lines 35-40, which discusses a certain number of bits describing the address of the main store and a certain number of bits describing the location in the cache.

The difference between instant claim and the explicit recitation of Brant et al is performing such storage in response to a power failure, although Brant et al do disclose a failure in general in column 2, line 14, for example.

However, Katz et al disclose storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information as a result of the power failure, as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

With respect to independent claim 14, receiving data access requests at the cache is disclosed in column 4, lines 32-41.

The cache having storage formed by at least two disks is disclosed in figure 1, #16, and discussed in the abstract, for example.

The requests specifying respective data storage addresses is disclosed in the title and abstract. A cache is, by definition, a memory subsystem in which frequently used values are duplicated for quick access. A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory. A cache is useful between a processor and a main memory when main memory accesses are slow compared to the processor speed, because the cache memory is always faster than the main memory.

As discussed in column 4, lines 9-19, the storage controller receives data access requests (via front-end interface #11) that specify respective data storage addresses, retrieves data identified by the data storage addresses, and services at least one of the requests received at the front-end interface using data stored in the at least two disks. The at least one of the requests serviced are, of course, the ones in which the cache holds the information for. As the definition of cache suggests, as discussed *supra*, if the

cache does not have the information, the cache does not service the request *per se*, but the main memory, or larger storage does (via back-end interface #24). This is also discussed in column 4, lines 9-19. As the data stored in the at least two disks that make up the cache is stored striped in a RAID configuration, as discussed in column 5, lines 34, 36, and 44, for example. This is significant as RAID 3, for example, as mentioned in column 5, line 44, adds redundant information in the form of parity to a parallel access striped array, by definition. Striping is assigning blocks of storage in regular sequence to all of an array's disks. Therefore, Brant et al uses stores the data in the at least two disks that make up the cache by striping it in a RAID 3 configuration, for example.

The cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy, as discussed in column 7, lines 35-40, which discusses a certain number of bits describing the address of the main store and a certain number of bits describing the location in the cache.

The difference between instant claim and the explicit recitation of Brant et al is performing such storage in response to a power failure, although Brant et al do disclose a failure in general in column 2, line 14, for example.

However, Katz et al disclose storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information as a result of the power failure, as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

With respect to independent claim 23, a data storage system is disclosed in figure 1.

A 'back-end' storage system is disclosed in figure 1, #25. Data storage spaces all have addresses, or locations for storing data – this is inherent, by definition of an address and by definition of data storage. The addresses identifying "blocks" of storage is also inherent – a block of storage may be any size, as applicants have not limited such size in the instant claim. Therefore, examiner interprets a "block" to be one memory location.

A cache for the back-end storage system is disclosed in figure 1 as #16 as well as in the title and abstract. A cache is, by definition, a memory subsystem in which frequently used values are duplicated for quick access. A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory. A cache is useful between a processor and a main memory when main memory accesses are slow compared to the processor speed, because the cache memory is always faster than the main memory.

A front-end interface is disclosed in figure 1, #11, as the "host interface" in column 6, line 17. Receiving I/O requests that specify respective addresses of back-

end storage blocks is disclosed in column 5, line 35, which refers to I/O rates. An I/O rate must be calculated from I/O requests. Therefore, I/O requests are inherent in a system with an I/O rate being measured. Also, it is important to note that I/O is "Input/Output," which is defined as the complementary tasks of gathering data for a computer or program to work with, and of making the results of the computer's activities available to the user or to other computer processes. This is exactly what figure 1 shows – data is gathered from the disk-based disk cache and supplied to the host.

A back-end interface is disclosed in figure 1 as #24, as the "interface to the mass data storage subsystem" in column 6, line 18.

Cache storage being formed by at least two disks is disclosed in figure 1, #16, and discussed in the abstract, for example. The disks having platter diameters [of] less than 3.5 inches is disclosed in column 3, lines 45-46.

A cache manager is disclosed as the storage controller in figure 1, #20. As discussed in column 4, lines 9-19, the storage controller receives data access requests (via front-end interface #11) that specify respective data storage addresses, retrieves data identified by the data storage addresses, and services at least one of the I/O requests received at the front-end interface using data stored in the cache storage. The at least one of the requests serviced are, of course, the ones in which the cache holds the information for. As the definition of cache suggests, as discussed *supra*, if the cache does not have the information, the cache does not service the request *per se*, but the main memory, or larger storage does (via back-end interface #24). This is also discussed in column 4, lines 9-19. As the data stored in the at least two disks that make

up the cache is stored striped in a RAID configuration, as discussed in column 5, lines 34, 36, and 44, for example. This is significant as RAID 3, for example, as mentioned in column 5, line 44, adds redundant information in the form of parity to a parallel access striped array, by definition. Striping is assigning blocks of storage in regular sequence to all of an array's disks. Therefore, Brant et al uses stores the data in the at least two disks that make up the cache by striping it in a RAID 3 configuration, for example.

The cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy, as discussed in column 7, lines 35-40, which discusses a certain number of bits describing the address of the main store and a certain number of bits describing the location in the cache.

The difference between instant claim and the explicit recitation of Brant et al is performing such storage in response to a power failure, although Brant et al do disclose a failure in general in column 2, line 14, for example.

However, Katz et al disclose storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information as a result of the power failure,

as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

With respect to claims 2 and 15, Brant et al disclose the front-end interface comprising an interface conforming to “a protocol” is disclosed in column 6, lines 39-41, for example, where the protocol is a “SCSI-type connection[s].”

With respect to claims 3 and 16, Brant et al disclose the protocol comprising at least one of SCSI, Fibre Channel, “INFINIBAND,” and IDE is disclosed in column 6, line 22, which identifies IDE, as well as line 41, which identifies SCSI and associates SCSI with the front-end interface (host interface #11).

With respect to claim 4, Brant et al disclose the disks comprising disks having platters less than 3.5 inches in diameter is disclosed in column 3, lines 44-46.

With respect to claim 5, the disks comprising disks having at least one of the following platter sizes: 2.5 inches, 1.8 inches, and 1 inch in diameter is disclosed in column 3, lines 44-46, which not only discloses the 1.8 inch diameter disk, but also states that “(or smaller)” [would work in the invention]. “Or smaller” would include the 1 inch diameter as well.

With respect to claims 6 and 19, Brant et al disclose the cache implementing a RAID scheme using the disks is disclosed in column 5, lines 34, 36, and 44, in general. RAID stands for, “Redundant Array of Independent Disks.” In this case, column 5, line 59 recites, “Controller 20 can include independent paths to write data to its memory in a mirrored fashion.” Mirroring is redundant storage of data. The cache being an Array is disclosed in column 4, line 15, for example. Figure 1 clearly shows separate disks, and,

therefore, independent disks. Therefore, RAID is explicitly disclosed embodied in the invention of Brant et al.

With respect to claim 7, Brant et al disclose the cache performing at least one of the following operations is disclosed below. Examiner reminds applicants that as the claim language stands, only one of the following limitations are required to be anticipated by the instant prior art of record. However, it happens that all of the following limitations are anticipated by the cited prior art of record as follows:

Requesting data from a back-end storage system (see column 6, lines 50-51);

Retrieving requested data from the [at least two] disks [making up the cache] (see column 4, lines 9-19);

Sending data to the back-end system for writing (column 6, lines 50-51);

Determining the location of back-end system data within the [at least two] disks [making up the cache] (column 4, lines 32-48).

Removing data from the [at least two] disks [making up the cache] (column 4, lines 42-44).

With respect to claim 8, Brant et al disclose the addresses specifying storage locations of a back-end storage system that includes a collection of one or more disks is disclosed in the definition of a cache, discussed *supra*, in that, "A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory." The back-end storage system is the main memory. The back-end storage system comprising a collection of one or more disks is disclosed in figure 1, #25 and further discussed in column 6, lines 57-58, for example.

With respect to claim 9, Brant et al disclose the requests comprising I/O requests is disclosed in column 5, line 35, which refers to I/O rates. An I/O rate must be calculated from I/O requests. Therefore, I/O requests are inherent in a system with an I/O rate being measured. Also, it is important to note that I/O is "Input/Output," which is defined as the complementary tasks of gathering data for a computer or program to work with, and of making the results of the computer's activities available to the user or to other computer processes. This is exactly what figure 1 shows – data is gathered from the disk-based disk cache and supplied to the host.

With respect to claims 10 and 21, Brant et al disclose the data storage addresses comprising data storage addresses within an address space is inherent by definition of an address, which merely denotes a location of memory in which something is, or may be, stored. The address space is the disk-based disk cache.

With respect to claims 11 and 22, Brant et al disclose the address space comprising an address space of back-end storage is disclosed by the definition of a cache, as discussed with respect to claims 1 and 8, for example.

With respect to claim 12, Brant et al disclose a storage hierarchy associated with the various contemporary storage configurations, in order of accessing speed, in column 5, lines 10-28. Column 5, lines 29-31, recites that, "A storage subsystem that has the MB cost of disk coupled with the performance of many disks operated in parallel can fill several intermediate slots in this hierarchy." By introducing the storage subsystem at another slot of the hierarchy that has a level below it of slower access storage, the subsystem effectively becomes another cache, and its respective address

space becomes another cache's address space. An example of another level of memory that would have a slower access time might be a tape library or the like.

With respect to claim 13, Brant et al disclose the cache storage having more than one disk spindle is inherent in a duplex mirrored disk subsystem, which is disclosed in column 3, line 1. A disk spindle is an axle for mounting a disk. A duplex disk subsystem is a system of two spindles, one of which is active while the other remains on standby, ready to take over processing if the active spindle malfunctions.

With respect to claim 17, Brant et al disclose the requests comprising at least one read request is disclosed in column 4, line 25, for example.

With respect to claim 18, Brant et al disclose servicing the requests comprising retrieving data from the back-end storage (main memory) and storing the data in at least one of the disks is inherent in the definition of caching. The addresses specifying storage locations of a back-end storage system that includes a collection of one or more disks is disclosed in the definition of a cache, discussed *supra*, in that, "A cache stores the contents of frequently accessed main memory locations and the addresses where these data items are stored in the main memory." The back-end storage system is the main memory. The back-end storage system comprising a collection of one or more disks is disclosed in figure 1, #25 and further discussed in column 6, lines 57-58, for example.

With respect to claim 20, Brant et al disclose servicing the requests comprising determining whether the collection of disks currently stores the requested data is disclosed in column 4, lines 39-41, for example.

With respect to claims 25 and 30, Brant et al disclose the cache further comprising at least one interface conforming to a protocol to allow at least one additional disk to be connected to said storage is disclosed in column 6, lines 39-41, for example, where the protocol is a SCSI-type connection(s). A SCSI interface is inherently used to connect CPU's or hosts to SCSI peripheral devices (note the plural), such as many hard disks and printers, using SCSI ports, which inherently provide logical connections between the computer and all of the devices on the SCSI bus. Therefore, inherently, a SCSI bus connection through SCSI ports allows at least one additional disk to be connected through the interface.

With respect to claim 38, the cache manager storing identification data identifying addresses within said cache storage where data is stored and the corresponding address at a back-end storage area where data is stored is inherent in a direct-mapped strategy, as discussed in column 7, lines 35-40, which discusses a certain number of bits describing the address of the main store and a certain number of bits describing the location in the cache.

The difference between instant claim and the explicit recitation of Brant et al is performing such storage in response to a power failure, although Brant et al do disclose a failure in general in column 2, line 14, for example.

However, Katz et al disclose storing a "journal of information concerning the write operations to be performed," which is interpreted by examiner to be the direct mapping information, as is elaborated in column 12, lines 35-50, for example) in column 12, lines 19-20, in response to a power failure, as disclosed in column 12, line 24.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of the copending application and Katz et al before him/her, to utilize the direct-mapping information of Brant et al specifically as a result of the power failure, as it enables recovery from that power failure, as disclosed by Katz et al in column 12, lines 28-30, so data is not lost.

With respect to claims 39-42, the combination of Brant et al and Katz et al disclose the limitations of the claims upon which the instant claims depend, as discussed supra.

The difference between said combination of Brant et al and Katz et al is the explicit recitation of the identification data corresponding to cache locations of "deferred writes." A "deferred write" is interpreted by examiner to be a write before the start of the write operation.

However, Katz et al disclose identification data corresponding to locations of deferred writes in column 12, lines 35-50, which discuss storing the identification data before the start of any write operation.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention, having the teachings of Brant et al and Katz et al before him/her, to store the identification data before the start of any write operation because it enables such data to be used to recover from a power failure, as disclosed in column 12, lines 29-31, for example.

Response to Arguments

With respect to applicants' argument that Brant et al does not anticipate each and every limitation in the amended claims, it is rendered moot by the new grounds for rejection presented *supra*.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Christian P. Chace whose telephone number is 703.306.5903. The examiner can normally be reached on 9-4-5.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on 703.308.1756. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).



Christian P. Chace